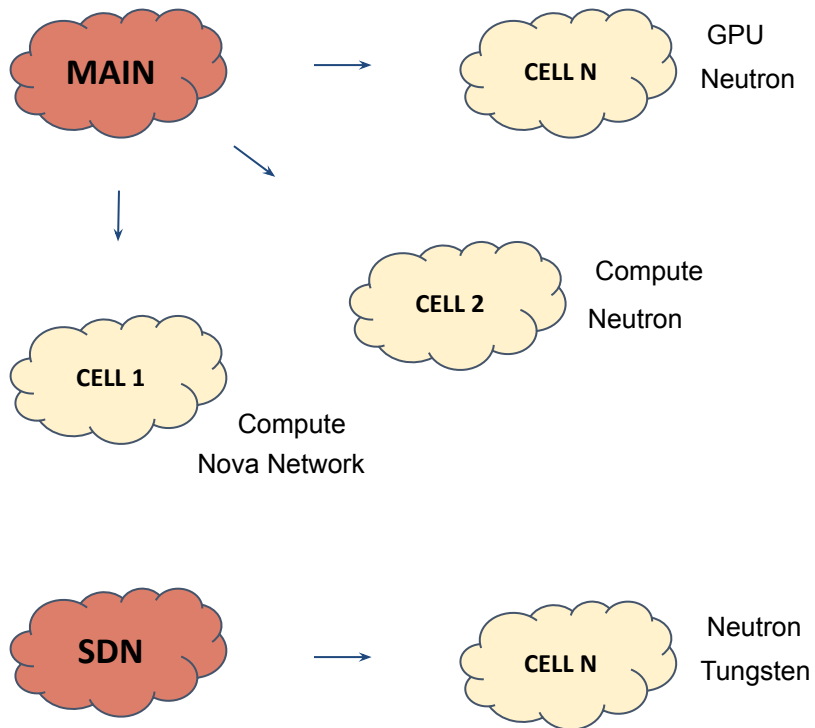


Networking at CERN

Ricardo Rocha

On behalf of the CERN Cloud team

State today...



Production setup

Multiple Regions

Mix of Nova Network, Neutron


SDN setup

Ongoing work, Tungsten

Networking at CERN

Flat, segmented network (broadcast domains)

Hierarchy of Primary (hypervisors) and Secondary (VMs) services



Networking at CERN

A device database as the single source of truth

Integration via a Neutron mechanism driver + extensions



Primary Services



Secondary Services



Hypervisors



Virtual Machines

IPv4

IPv6

DNS

Aliases

Aliases

IPv6 Readiness

Ownership

...

Neutron at CERN

Linuxbridge, ML2

Extension to manage *Clusters*: similar to Neutron Segments

```
neutron cluster-list
```

id	name	subnets
...	VMPOOL SXXXX-C-IPZZZ	... 188.xxx.yy.zz/22
...	VMPOOL SBBBB-C-IPWWW	... 137.aaa.bb.ccc/25
		... 137.bbb.cc.0/25
		... 137.bbb.dd.0/25

Neutron at CERN

Linuxbridge, ML2

Extension to manage *Clusters*: similar to Neutron Segments

Extension to manage *Hypervisor / Subnet* mappings

```
neutron host p06253927y321a1
```

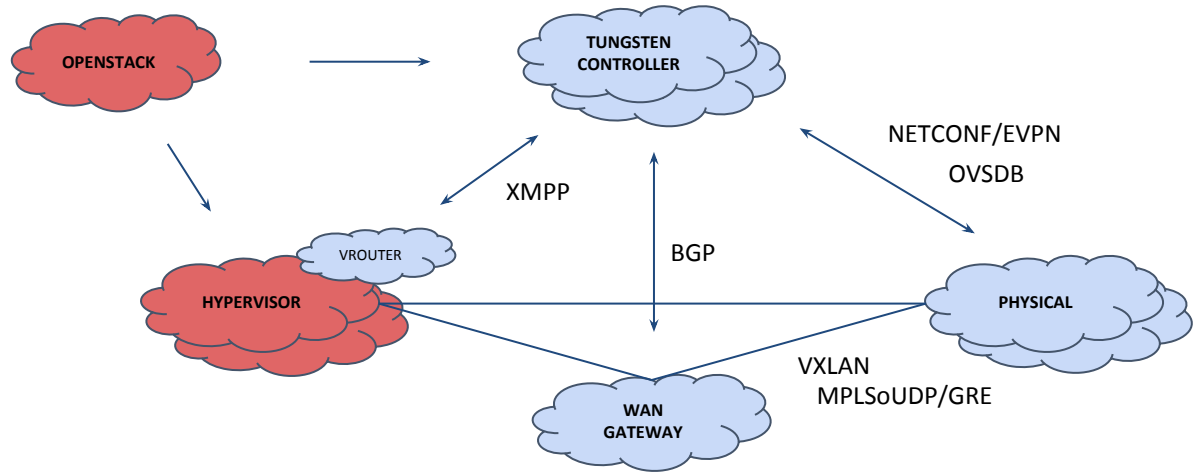
Field	Value
all_subnets	4ca09148-32b5-4da4-95f9-35e83e2e1984
available_random_subnet	4ca09148-32b5-4da4-95f9-35e83e2e1984
available_subnets	4ca09148-32b5-4da4-95f9-35e83e2e1984
least_available_subnet	4ca09148-32b5-4da4-95f9-35e83e2e1984
most_available_subnet	4ca09148-32b5-4da4-95f9-35e83e2e1984

An exciting new world coming next...


New region with software defined networking

Floating IPs, LBaaS, Overlay Networks, QoS?

Only the API is left of Neutron, everything delegated to Tungsten Fabric



An exciting new world coming next...



Split infrastructure as we scale and gain confidence with the new deployment

But the underlay network is still the same, routable

First up: LBaaSv2 instances running in Region *sdn*, members anywhere

	Neutron/OpenVSwitch	OpenDaylight	OVN
DHCP	Neutron	Neutron/Built-in	Built-in
Floating IPs	Yes	Yes	Yes
Distributed Routing	Only with DVR	Yes	Yes
Tunneling Protocols	vxlan / GRE / geneve	vxlan / GRE / geneve	vxlan / geneve
Security Groups	IPTables	OpenFlow Native	OpenFlow Native + Logging
Load Balancing	Octavia	Octavia	Octavia / OVN Native
Acceleration	Limited DPDK	DPDK	DPDK
Tracing	tcpdump	tcpdump	ovn-trace
Physical Switch Integr.	L2 / L3	L2 / L3	L2 / L3